CSSE4004/7014
The University of Queensland

Tutorial 8

Consistency and replication

1. Explain in your own words what the main reason is for actually considering weak consistency.

2. Consider each of the models of distributed shared memory in (a)-(b) and explain how you would classify them according to the models of consistency and replication in the lecture. Will any have consistency problems?

    a. All processors keep a counter representing global time. Every time a write occurs, all processors have to be signalled, and stalled, until the write is complete. A copy is broadcast to all caches, along with the clock counter of the writer, and replaces any existing copy. If it is not currently cached, the new value is discarded if the processor's local clock counter is ahead of the broadcast time. If it's not, it buffers the new value until its clock has caught up, at which point the new value is cached.

    b. Before a write can occur, the memory containing the data to be written has to be flushed from memory of all other processors in the system. Any subsequent reads will result in a miss, which has to be serviced by modifying the written copy to sharable, and copying it over the network to the reader.

3. How do you think a large free email service with a worldwide presence, such as *yahoo* operates? Assume that replication to a local server is a useful strategy to minimize latency, but that a fair fraction of users are highly mobile. Consider a range of potential consistency and replication models, and explain which you think best, and which you consider to be unsuitable. Consider possibilities of both *client-centric* and *data-centric* solutions.

4. You plan to build a distributed system in Java, and would like to maximize leverage from standard Java APIs. What extensions to standard Java APIs would you need to support replication?

5. How could you use read-your-writes consistency to ensure that you see the latest copy of a web page you've just updated yourself? What would you have to add to a standard web browser and cache on a personal computer?

6. A file is replicated on 6 servers.
    a. List all the combinations of read and write quorums permitted by Gifford's scheme.
    b. Do all of these combinations make sense?
    c. What is the significance of a read quorum of 1?

7. What kind of consistency would you use to implement a web-based flight booking system, with multiple airlines, accessible to users around the world? Could more than one apply to different parts of the system?

8. For each of the following, explain why the example does not meet the specific consistency requirements:
    a. Figure (b) strict consistency on slide 11.
    b. Figure (b) sequential consistency on slide 12.
    c. Figure (a) causal consistency on slide 15.
    d. Figure (b) monotonic reads on slide 28.
    e. Figure (b) monotonic writes on slide 30.
    f. Figure (b) read-your-writes on slide 31.

9. Explain general issues in choosing between client-centric and data-centric consistency models.